

LA-UR-21-26141

Approved for public release; distribution is unlimited.

Title: Data and Insights for the KRUSTY Nuclear-Powered Tests

Author(s): Poston, David Irvin

Intended for: Report

Issued: 2021-06-29

Disclaimer:

Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by Triad National Security, LLC for the National Nuclear Security Administration of U.S. Department of Energy under contract 89233218CNA000001. By approving this article, the publisher recognizes that the U.S. Government retains nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy. Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness.

DATA AND INSIGHTS FOR THE KRUSTY NUCLEAR-POWERED TESTS

David Poston, Los Alamos National Laboratory, poston@lanl.gov

Note: I have written this paper to document everything I can think of that might be of interest regarding the KRUSTY tests, in particular to aid in benchmarking. The document is rather informally written, and I am writing it as a sole author because there is a lot of opinion incorporated, and neither I nor the team members have adequate time to fact check and compile everything, or come to a consensus on opinions. Of course the body of work can be attributed to numerous people associated with NASA and the NNSA, most notably Marc Gibson, Patrick McClure, Tom Godfroy and Rene Sanchez. Note: access to the raw data files may be limited, and at the time of this publication a permanent repository has not been established. A later revision of this document will contain more information about accessing the data (as listed in Figure 1).

INTRODUCTION

This document, and the associated files, describe all of the data from the KRUSTY warm criticals and the KRUSTY nuclear system test. The warm criticals were a group of tests where fission power heated the core and the feedback response was studied. There were 3 experiments, the 15-cent run (which was a true “free run”), followed by 30-cent and 60-cent runs, where reactivity was added to heat the core to higher temperatures. The nuclear system test (often referred to as the full-run or the final-run), engaged the power conversion system and evaluated system performance. These experiments are described in a pair of NUCLEAR TECHNOLOGY papers^{1,2}, which includes plots of the data that this document-set includes. Note that a lot of information in this document is based on my interpretation and could be wrong, but I’ve included it to potentially save some people a lot of time and headaches by trying to figure out which instrumentation to believe, but beware treating anything I say as fact. I’ve also included several post-irradiation photos, to show some thermocouple locations and indicate possible changes in emissivity and/or geometry (although no changes in geometry were indicated).

Use of this Data for “Warm” Criticality Benchmarking

The warm criticals might not be technically considered “criticals”, because they did not produce “known” values of criticality/reactivity at any specific temperature or configuration (other than they all started with a 15-cent free run at cold startup). The 30-cent and 60-cent criticals moved the platen (lift table) to a point where we had guessed 30 or 60 cents were to be inserted, but only the ~15 cent insertions were “known”. The use of the term “known” refers to a value of reactivity that can be simply and accurately deduced; e.g. via the logarithmic slope of power (neutron flux) in a zero-power cold critical, or if a reactor is operating a steady state (and thus has reactivity near zero, apart from the contribution from the source). It may be possible to estimate a change in reactivity during the KRUSTY runs by calculating the change in slope of power, but 2nd order effects will add uncertainty to this estimate (depending on the specifics of the transient at any point in time). Also, temperature diagnostics are uncertain and limited to specific areas, and thermocouple lag will make it hard to match up reactivity and temperature on a temporal basis. The best criticality benchmark from this data, other than the 15-cent free run, would probably be near the end of the 60-cent run; reactivity is close to zero, thermocouples have all well-thermally-coupled with almost no time-lag, and thermal gradients within the fuel are close to zero (so there is less uncertainty on the difference between inner and outer fuel temperature). Although, as mentioned previously, there is no way to accurately know if 60 cents was actually inserted.

Use of this Data for Transient/Dynamic Benchmarking

One of the purposes of the warm criticals was to provide data for dynamic or transient benchmarking. The warm criticals were actually used as part of the safety approval process to gain permission to proceed with the KRUSTY full-run. The 15-cent and 30-cent runs provided data to benchmark the transient code FRINK, so that it could hopefully predict the results of the 60-cent run (which it ultimately did). The experimental data is now of great importance because there are numerous efforts to create “multi-physics” codes that can model reactor behavior, and this data will provide simple transients for benchmarking those codes. The best thing about KRUSTY for transient benchmarking is that it is effectively a point-kinetics reactor, which will allow simpler methods and/or less rigorous modeling detail to attempt a benchmark.

The warm criticals offer nice benchmark cases because the power losses from the core are passive; i.e. the complexities of the KRUSTY power conversion system do not need to be modeled. For all of the warm criticals, the vast majority of power loss is through the multi-layer insulation (MLI); although for the 60-cent critical the temperatures got warm enough to get a little heat transfer along the heat pipes, which might slightly impact benchmarking. In these cases, the MLI conductance is by far the biggest uncertainty, and for most benchmarking efforts I assume it will be the key “fudge factor” in trying to match the data. The heat transfer through the radial wraps of MLI are hard enough to predict with a physics-based model, but even worse is the axial MLI, which is under compression (the lower axial MLI likely experienced significant temperature dependent compression due to structural thermal expansion). Each of the warm criticals has a cool-down period after scram which will be very valuable to benchmark/fudge the MLI conductance as a function of temperature. The other most significant fudge-factors are the effective thermal conductance between the core, heat pipes, and core ring-clamps. The thermal conductance appears to be very good (close to ideal) at high temperature, but unfortunately not as good at low temperature (room temperature 100 C, 200 C?). The design transient code for KRUSTY, FRINK, was benchmarked to the 15-cent and 30-cent runs by varying these fudge factors to match each “region” of each transient, and was then able to predict the 60-cent run very well. A good benchmarking effort should be able to create a model that mimics the results of all three warm criticals using a consistent set of assumptions.

Benchmarking for the nuclear system test, or full-run, introduces a lot of complicating factors. Most difficult are the startup of the heat pipes and heat draw/loss from the Stirling simulators. Some of these difficulties are discussed in Reference #2.

KRUSTY DESCRIPTION

Almost everything needed to “adequately” benchmark the KRUSTY nuclear-powered testing can be found in previous publications, most notably NUCLEAR TECHNOLOGY papers^{1,2,3,4,5,6,7}. Materials, densities, and dimensions are provided to generate an adequately detailed model, i.e. the lack of additional detail should be of little concern versus the large uncertainties in contact resistance, MLI conductance, and the experimental measurements themselves. The dimensions of unspecified components can be estimated to an adequate level based on drawings/photos, and standard material properties can be found or estimated for density, conductivity, specific heat, expansion coefficients, and emissivities.

For the “reactor” an extremely detailed description of most reactor components can be found in a criticality benchmarking paper by Smith⁸. This is not the exact KRUSTY configuration, but the dimensions, density, and material specs for all key components are provided.

The reason I am writing this document is to at least provide my understanding of every issue I can think of, without the benefit of a new project track down every detail.

The level of detail required in specific areas is certainly different if one is creating a criticality benchmark as opposed to a dynamic benchmark. The criticality benchmark will be focused on material specs, detailed dimensions, and bulk temperatures, whereas the biggest concern in creating a dynamic benchmark is generally heat transfer characteristics. In either case, a good criticality model is needed to create reactivity feedback curves or feedback coefficients, except in cases where the transient tool uses a neutronics model as part of the integrated solution.

One of the biggest uncertainties is the MLI. The MLI basic dimensions are given in the design paper³, but to clarify, the MLI came prepackaged as a “wrap” of 5-layers of .001” Mo separated by 4-layers of quartz thread. The thread was .004” Astroquartz (style 507, finish 9779). In this and other documents, this “wrap” is referred to as 4 layers of MLI, even though there are 5 layers of Mo. This is why the MLI is always in factors of 4; 4 layers between axial reflectors and fuel = 1 wrap, 8 layers in the radial MLI surrounding the core = 2 wraps, and 16 layers around the heat pipes = 4 wraps. If there are multiple super-wraps (wraps of 4-layer wraps), the Mo-Mo interface is not referred to as an MLI layer because there is no thread providing separation. Any reference to MLI that refers to 5 or multiples of 5 is likely just labeled differently, but would still be this same 5-Mo, 4-thread configuration.

The biggest lacking description are the dimensions and masses of the Stirling converters and thermal simulators. I have found that including the thermal inertia of the simulators definitely creates a better benchmark, although I’ve had to estimate the masses. A schematic and photo of the convertors is shown in Reference #4. The simulators have the same hot-end dimensions with a similar copper heat acceptor inside. Within the simulator, nitrogen gas flows down an outer annulus to the base of the simulator, reverses path and flows over the heat acceptor (a planar heat-exchanger of thin copper). The gas outlet thermocouple is perhaps 1” downstream of the acceptor. A double wall exists between the down-flow annulus and the up-flow tube, which contains 4-layers of MLI (i.e. 1 wrap). Perhaps NASA might publish more detail on the simulators in the future.

DATA FILE DESCRIPTIONS

File listing

There are 3 types of files that contain the raw data. The neutron detector readings are in the Comet.dat files, the comet platen position is in Platen.dbf files, and all other data of interest is in the NASA.txt files. Figure 1 below shows a screen capture of the file listing of all files in the database.


















Name	Date modified	Type	Size
 Comet_07Mar2018_15cent_run.dat	3/7/2018 5:37 PM	DAT File	910 KB
 Comet_08Mar2018_30cent_run.dat	3/9/2018 1:45 AM	DAT File	1,034 KB
 Comet_14Mar2018_60cent_run.dat	3/15/2018 4:38 PM	DAT File	1,575 KB
 Comet_20Mar2018_FullRun_Shift1.dat	3/23/2018 1:42 PM	DAT File	1,519 KB
 Comet_20Mar2018_FullRun_Shift2.dat	3/23/2018 1:42 PM	DAT File	1,544 KB
 Comet_21Mar2018_FullRun_Shift3.dat	3/23/2018 1:42 PM	DAT File	1,191 KB
 Comet_21Mar2018_FullRun_Shift4.dat	3/23/2018 1:43 PM	DAT File	1,375 KB
 Comet_21Mar2018_FullRun_Shift5.dat	3/23/2018 1:42 PM	DAT File	648 KB
 NASA_07Mar2018_15cent_run.txt	3/7/2018 2:11 PM	Text Document	28,582 KB
 NASA_08Mar2018_30cent_run.txt	3/8/2018 3:01 PM	Text Document	40,876 KB
 NASA_14Mar2018_60cent_run.txt	3/18/2018 4:46 PM	Text Document	160,831 KB
 NASA_20Mar2018_FullRun_Day0.txt	3/21/2018 2:11 PM	Text Document	215,285 KB
 NASA_21Mar2018_FullRun_Day1.txt	3/29/2018 4:18 PM	Text Document	174,776 KB
 NASA_22Mar2018_FullRun_Day2.txt	3/29/2018 4:18 PM	Text Document	165,434 KB
 NASA_23Mar2018_FullRun_Day3.txt	3/29/2018 4:18 PM	Text Document	162,223 KB
 Platen_20Mar2019_Fullrun0.dbf	3/23/2018 1:42 PM	DBF File	8,816 KB
 Platen_21Mar2018_Fullrun1.dbf	3/23/2018 1:42 PM	DBF File	8,756 KB

Figure 1. Files containing the raw KRUSTY data.

NASA.txt Files

These files are cumbersome and somewhat confusing. First, there are lot of data columns that were originally used for preliminary and electrically-heated testing, and ended up being irrelevant for KRUSTY. Also, some of the columns are calculated values (e.g. meant to inform the Stirling simulator flow setting), and thus are not experimental data. In addition, other complications led to a rather confusing layout of the NASA data in the .txt tiles.

Initially, the column-wise order of the data was logical, and was organized according to how the TC wires tied in to the NASA Rack. Two things led to a rather disorganized output file. When at NCERC some of the data was routed to bypass the NASA rack (which sat in hallway rather close to KRUSTY) and instead go directly to the control room via dedicated channels in the facility – these are labeled as NSV channels. Previously during the DUFF testing we lost our connection to the data computer during the test, so we wanted to make sure that at least some thermocouple data would make it into the control room if the computer/rack tripped. As a result, those thermocouple wires were routed to a different physical location, resulting in a different column position in the txt file. The second cause of column confusion was that we indeed had a loss of data in the NASA rack during the 30-cent run. As a result, we moved many more of the thermocouples to the NSV system, to make sure we got the most important data if the NASA rack computer crashed again. Thus, several TC readings changed column between the 30-cent and 60-cent run, and the labeling at the top of some columns is messed up. The good news is that we fortified the shielding of the NASA rack before the 60-cent run and the computer trip did not occur again (during the full run either). There is a long story behind how we ended up with a rather poorly shielded

computer; we initially had it well-shielded, but several changes (caused by other requirements) all added up in the wrong direction towards more dose to the NASA rack. This was probably our biggest embarrassment/failure, since we had supposedly learned our lesson about single-event upsets during the DUFF test. However, at least we learned the lesson well enough to send some signals the NSV path, instead of relying fully on the rack computer – and it was only the 30-cent run where we lost any data.

Comet.dat Files

These files contain the readings of the 3 facility neutron detectors mounted on the room walls surrounding KRUSTY. The neutron detectors are linear ^3He detectors, for which the readout is amps. The readings from each detector stayed in good proportion throughout the testing, which makes sense because there should be no azimuthal dependence on platen position (except for 3rd order shine effects). The average of the 3 readings could be used, but I like to use Linear1, because both Linear2 and Linear3 have a small amount of noise/variability in the readings. As mentioned in the results paper, this Amp reading has to be normalized to fission power, which can be approximated by integrating area under the neutron flux curve, and normalizing it to Joules by using the thermal capacitance ($C_p \cdot \text{mass}$) of the core. Actually, immediately after the 15-cent run I ran the FRINK system model (which included passive losses) and determined the normalization that best matched the transient. Then again after the 30 cent, which resulted in a <1% change, and the used that normalization to predict the 60 cent run; for the latter runs, both models provided the same normalization. The normalization I came up with for fission power is $\sim 5.5 \times 10^8$ Watts/Amp when using the data from Linear1. The change in magnitude of each detector will be a small function of the platen position (e.g. change in reflector gap and the streaming path at the lower shield), which might have small effect on the normalization of detector reading and fission power during the full run, but should be minor.

There were also fission foils used to estimate flux/fission rate, which provided corroborating results, but the uncertainty of that technique is likely larger than the simple approximation based on thermal capacitance. And the full system model normalization should be the best, assuming the model is good of course. This results in a potential modeling catch-22, but if the normalization is significantly off, the model for all 4 runs should not match very well.

Platen.dbf Files

These files contain the platen position during the full run, specifically the platen gap size in inches, where 0.00" would represent fully closed and the nominal fully withdrawn condition is 25.00". This data was not recorded for the warm criticals, but the initial and final platen positions were recorded in the log book. Also, the data for the full-run is only printed to the 0.01" inch, whereas the log-book had specific positions to the 0.001" (and KRUSTY is sensitive enough that every mil is significant). The table below shows the platen positions for the warm criticals, and a few "3-digit" values for the full run. Note that all nuclear-powered cases had 1.0" of BeO in the shim pan.

Table 1: Platen loading and gap size for nuclear-powered tests.

Test	Platen BeO Height	Condition	Platen gap
15-cent run	11.250" inner, 11.125" outer	Initial ramp (~15.7 cents)	0.205"
30-cent run	11.250" inner, 11.125" outer	Initial ramp (~15.5 cents)	0.208"
30-cent run		Steady platen position	0.157"

60-cent run	11.250" inner, 11.125" outer	Initial ramp (~15.7 cents)	0.209"
60-cent run		Steady platen position	0.056"
full-run	11.625" inner, 11.625" outer	Initial ramp (~15.5 cents)	0.665"
full-run		T=~1.0, pause for power rise	0.315"
full-run		T=~1.5, end insertion	0.253"
full-run		T=~6.5, 1 st reactivity adjust	0.256"
full-run		T=~6.8, 2 nd reactivity adjust	0.265"
full-run		T=~7.1, 3 rd reactivity adjust	0.275"
full-run		T=~16.0, lower/remove transient	0.297"
full-run		T=~17.0, lower/remove transient	0.317"
full-run		T=~18.0, raise/add transient	0.257"
full-run		T=~19.0 lower, final position	0.273"

As noted, all of the platen positions for the full-run are in the platen.dbf file, but this table adds another significant digit for a more accurate benchmark at various times.

Time synchronization:

The time stamp information on each of the data files is different. The NASA.txt files read PST, while the platen.dbf files read PDT. The change to Daylight Savings Time occurred the weekend before the full run, but the NASA computers did not automatically adjust from PST to PDT. The difference was noticed, but it was decided not to mess with the NASA computers and proceed as is (i.e. it was deemed not worth risking any sort of reboot or OS issue on these old systems). The Comet.dat files are not strictly time stamped. The system is hardwired for a time limit of 8 hours of data (the machine had never been used for extended operations before). The starting time of the data collection is given at the top of the file, but only printed to the minute (not second). Thus, this data must be lined-up "by hand", by looking at the results. This is very hard to do at the start of the transient, but easy to line up at scram, matching the drop in neutron current with the drop in temperatures. There was some debate as to whether there was some drift between the NASA and NCERC clocks over time during the testing campaign, but I think only a few seconds per day, and I think that the computers were synchronized often. If there was a little drift during the full run, it could be revealed with careful examination of the full run, by lining up transients when the platen is moved. Also, note that there is downtime between the Comet.dat files. Because of the 8 hours limit, someone would have to go to the computer, download the data, and then reset the system to begin recording again. This took varied amounts of time, but generally the dead-period (with no data) became shorter the more times the reset was performed.

THERMOCOUPLE INFORMATION

Fuel Thermocouples

There were 18 fuel TCs placed on the outer surface of the fuel. They were spring loaded by a "clip" to be in forceful contact with a .001" Mo foil wrap that surrounded the fuel. The reason for the Mo wrap and the decision not to weld the TCs in place was to protect to the fuel from damage – for both KRUSTY and later potential uses for the fuel. The edges of the rectangular spring mechanisms for some of the

TCs can be seen in Figure 2, tucked beneath the clamping rings. The TCs were placed in 3 axial locations (top, mid, bot) at 6 azimuthal locations (labeled as 0, 45, 135, 180, 225, 315 degrees).




Figure 2. Core photos

The “top” TCs are located under clamping ring #1, “mid” are under ring #3, and “bot” are ring #5; where ring #1 (i.e. the first that was installed) is the top ring. Also note that the azimuthal position is actually 22.5° to the left of the labeled position (22.5° clockwise as viewed from the top). The photo is taken head on at the 315 heat pipe and shows the 315 TCs (labeled as such because the wires travel along the 315° heat pipe), although the TCs are actually located half-way between the 315° and 0° HPs, thus are really at 337.5° .

It is important to note that these fuel TCs are essentially reading the average temperature of the Mo core wrap and the spring-clip itself. The Mo core wrap temperature should be very close to the fuel if a good spring force is applied, but the spring-clip will read a bit colder since it is in contact with the core clamping ring and can also radiate to the radial MLI. I don’t think the magnitude of power deposition in the sheath and wire itself should not be a major factor in the TC reading, but it might cause them to read a degree hotter than they would otherwise, depending on the power level. A model could be created to determine the impact of internal power deposition, however none of the experimental data indicate that it had a significant effect (e.g. there was no apparent “additional” drop in TC readings after the scram of the full-power run).

Also, the core was electrically heated to $\sim 800^\circ\text{C}$ prior to the warm criticals – to check out the power conversion system, and also out-gas prior to the nuclear-powered runs. The drawback of heating to

>800 C prior to the nuclear test is that the TC spring-clips were full compressed, and likely did not fully relax back to create good thermal contact near room temperature. This is evident in the TC data during warm-up and at lower temperatures, although some TCs were apparently had better thermal coupling than others. Note: I have come closest to matching the data from the 4 nuclear tests (15c, 30c, 60c, and full run) by creating an empirical correlation that goes from poor conductance to good conductance as a function of temperature; most important is the change from room temp to ~200 C.

Table 2 shows the name and location of the core TCs.

Table 2: Fuel Thermocouples

TC name	Axial elevation from fuel bottom (cm)	Azimuthal location from 0° heat pipe	15c/30c .txt file column	60c/final .txt file column	Full data from 30c	Apparent? Thermal coupling
c_1_0	6.25	22.5°	BO	IL		Average
c_2_0	14.58	22.5°	BP	HX		Good
c_3_0	18.75	22.5°	BQ	BQ		Average
c_1_45	6.25	67.5°	CF	CF		Average
c_2_45	14.58	67.5°	HW	HW	Yes	Average
c_3_45	22.92	67.5°	CH	CH		Poor
c_1_135	6.25	157.5°	DL	DL		Average
c_2_135	14.58	157.5°	IF	IF	Yes	Average
c_3_135	22.92	157.5°	DN	DN		Average
c_1_180	6.25	202.5°				Failed
c_2_180	14.58	202.5°				Failed
c_3_180	22.92	202.5°	EC	IG		Very Poor
c_1_225	6.25	247.5°	ER	ER		Average
c_2_225	14.58	247.5°	IK	IK	Yes	Average
c_3_225	22.92	247.5°	ET	IT		Very Poor
c_1_315	6.25	337.5°	FX	FX		Average
c_2_315	14.58	337.5°	FY	II		Average
c_3_315	22.92	337.5°				Failed

The “.txt file column” is the column the data is parsed into if the .txt file is loaded into Excel, or counting alphabetically column-wise from left to right. The “full data from 30c” column are the TCs that were routed through the NSV system, so data was recorded throughout the entire 30-cent run. The other TCs flat-lined when the power apparently reached ~2 kW, and did not return until ~3hours later.

Note that the c_3_0 TC was not under the expected clamping ring (highlighted in gold on the Table 1). I actually just noticed this when pasting in the photos above, which I wish I had noticed earlier because I had spent a lot of time trying to rationalize why c_3_0 read significantly higher than the other “top” TCs. Hopefully there are no other TCs not in the expected position, and careful inspection of other photos seems to indicate this is the only one (although not all angles were photographed).

The “apparent thermal coupling” is based on my observations, and thus based on opinion, not fact. Interpreting the “top” TCs is the trickiest, because that is the lowest temperature fuel region (due to lower power density), thus it takes longer for the core to expand and create good thermal contact via the spring-clip. “Average” coupling means that it appeared to me that the reading was +/- a few °C from their “peer” TCs. TC c_2_0 is rated “good” because it was the only TC that responded rather quickly to the transients, and read the highest; although “good” does not imply that it is perfectly bonded. Actually, my observation has been that the TCs at 0° and 335° were on average significantly higher than their azimuthal peers. There is no expected azimuthal variation during the warm crits (only in the final system test); however, it is possible that the core expanded/contracted in such a way during electrical heated system checkout that it left a higher spring-clip compression force on that side of the core. Likewise the only working 180° TC appears to have the worst coupling of all. So I would guess that the 0° and 335° TCs are reading closer to the actual fuel surface temperature than the others.

The TC c_2_0 was apparently the best coupled TC, but it also happened to “fail” and not provide readings during the 60 cent run (this failure is probably not related to the status of it being best coupled, but maybe a short or gap developed?). Then, about 2 hours into the full run, the TC started suddenly reading again, and back to the highest levels of any TC, indicating again that it might be the best coupled.

Reflector, Vessel, and Shield Thermocouples

These TCs are crucial to any benchmarking effort, because they indicate power loss from the core, combined with local power deposition. Unfortunately the number of TCs was limited by physical constraints and some of the readings have considerable uncertainty. Table 2 shows the name and relative location of the ex-core reactor TCs.

Table 3: Ex-Fuel Reactor Thermocouples

TC name	Location	.txt file column
ab_90	Inserted (unbonded) between axial BeO and MLI at fuel bottom, at 90°	EN
at_90	Inserted (unbonded) between axial BeO and MLI at fuel top, at 90°	EO
ab_270	Inserted (unbonded) between axial BeO and MLI at fuel bottom, at 270°	GK
at_270	Inserted (unbonded) between axial BeO and MLI at fuel top, at 270°	GL
1inch_45	Taped in slot of outer BeO ring, ~1” from ID, 1” above BeO bottom, at 45°	GM
1inch_225	Taped in slot of outer BeO ring, ~1” from ID, 1” above BeO bottom, at 225°	GN
6inch_45	Taped in slot of outer BeO ring, ~1” from ID, 6” above BeO bottom, at 45°	GO
6inch_225	Taped in slot of outer BeO ring, ~1” from ID, 6” above BeO bottom, at 225°	GP
Planten_45	Taped to bottom of platen (naming typo) about 5” from radial C/L, at 45°	GQ
Planten_225	Taped to bottom of platen about 5” from radial C/L, at 225°	GR
shield_0	Taped on OD of radial shield at axial midpoint, at 0°	GS
shield_45	Taped on OD of radial shield at axial midpoint, at 45°	GT
mc_90	Taped on OD of radial MLI (between MLI and vessel), axial center at 90°	EK
mc_270	Taped on OD of radial MLI (between MLI and vessel), axial center at 270°	FZ

The yellow tape that holds various TCs in place can be seen in place can be seen in Figure 3.




Figure 3. Ex-Core TC photos

There is no photo of the axial MLI TC positions, as they were slid in between the foil and BeO from either side. The TC junctions were likely to be 1" to 1.5" in from the OD. The platen TC location of 5" from C/L is a guess, but is not important because there should be relatively little radial variation in platen temperature.

Modeling of heat transfer to the lower axial reflector is complicated because downward axial expansion of the core compresses the core into the axial reflector, with the MLI sandwiched between. The axial reflector is supported by the bottom of the core vacuum can, which does not heat up much as the core, and therefore does not expand downward as much. The bigger question is where this compressing force is applied. The lower axial reflector consists of inner and outer pieces, and it is likely that they did not stack up to exactly the same height, so the compression may have occurred primarily between the fuel and inner BeO plug piece, or the fuel and the outer BeO annular piece. This might make a significant difference in how heat is lost from the fuel through the lower axial reflector, because there is a gap between the inner and outer BeO pieces. If the fuel compresses with the inner BeO piece, the heat might stream down the center to the boss on the vacuum-can lower cap. If the fuel compresses with the outer/annular BeO piece it might then radiate from the outer BeO piece to the radial portion of the vacuum-can (there is no radial MLI between the axial reflector and the vacuum-can).

Heat Pipe Thermocouples

The heat pipe TC data is pretty straightforward. There were 4 TCs per heat pipe, for a total of 24, and were placed under a thick wrap (16 layers) of MLI. Each of the TCs was welded to the HP wall, so the thermal coupling was good, although there were still several TCs that displayed questionable results.

TC name	.txt file column	Trusty?
hp_1_0	BR	Bad
hp_2_0	BS	Solid
hp_3_0	BT	Bad
hp_4_0	BU	Solid
hp_1_45	CI	Solid
hp_2_45	CJ	Solid
hp_3_45		Gone
hp_4_45		Gone
hp_1_90	CX	Bad
hp_2_90	CY	Solid
hp_3_90	CZ	Squirrelly
hp_4_90	DA	Squirrelly
hp_1_135	DO	Solid
hp_2_135	DP	Solid
hp_3_135	DQ	Solid
hp_4_135	DR	Solid
hp_1_180	ED	Solid
hp_2_180	EE	Solid
hp_3_180	EF	Solid
hp_4_180	EG	Solid
hp_1_225	EU	Squirrelly
hp_2_225	EV	Bad
hp_3_225	EW	Solid
hp_4_225	EX	Solid
hp_1_270	FJ	Squirrelly
hp_2_270	FK	Solid
hp_3_270	FL	Solid
hp_4_270	FM	Bad
hp_1_315	GA	Bad
hp_2_315	GB	Solid
hp_3_315	GC	Solid
hp_4_315	GD	Solid

In the table above, “solid” TCs read within a few degrees of their peer TCs. A “bad” TC clearly had issues. The “squirrely” TCs all had some strange behavior. TC hp_1_225 drifted higher from T=3 hr to T=6 hr of the full-run, and then stayed consistently higher. The hp_3_90 and hp_4_90 TCs seemed to take turns providing good data and then step jumping to apparently bogus values; hp_3_90 was sketchy up until T=13 hr, then looks good, while hp_4_90 looks good all the way to T=20 hr then jumps very high. Note that hp_90 is the heat pipe attached to a Stirling convertor (not simulator), and most of the changes in those TC readings begin when an engine is started or stopped. Vibration could play a role in calibration or signal, or it is not out of the question that the readings are true; i.e. the results are physical, in that one side of the condenser is hot early on (not seeing down-flow of Na), both are fine during most of the run, and the other side is hot late. TC hp_1_270 is also squirrely and attached to an engine, although it is down near the core and reads significantly lower than its peers early and late during the system test. Alternatively, the other TCs attached to engine heat pipes, hp_1_90, hp_2_90, hp_2_270 and hp_3_270, all appear to be solid.

The relative positioning of the TCs for the heat pipes and simulators are shown in Figure 4. In this photo, only hp_1 and hp_2 are already under MLI, but they are all under MLI in the final configuration.




Figure 4. Heat pipe and simulator TC locations

None of these TCs were routed through the NSV data path, and thus are not useful for most of the 30-cent run; regardless, they are mostly irrelevant for the 30-cent run except for perhaps a small amount of conduction to the hp_1 TCs.

Collector and Hot-End Thermocouples

These TCs were welded to the just above the heat pipe condenser interface, to the nickel heat collector of the Stirling convertors (labelled as e_hotend) and the simulators (s_collector), and also covered by 16-layers of MLI. The collectors can be seen in Figures 9 and 10a of Reference #4.

There were 3 TCs per convertor and 2 TCs per simulator. These TCs locations are shown for the simulators on Figure 4 above, as well as some of the later figures.

TC name	.txt file column	Trusty? based on full run data
s_1_0	BV	Solid
s_2_0	BW	Solid
s_1_45	CM	Solid
s_2_45	CN	Solid
e1_1_90	HZ	Solid
e1_2_90	CV	Solid
e1_3_90	CW	Solid
s_1_135	DS	Bad
s_2_135	DT	Solid
s_1_180	EH	Solid
s_2_180	EI	Solid
s_1_225	EY	Solid
s_2_225	EZ	Solid
e2_1_270	IN	Solid
e2_2_270	FH	Solid
e2_3_270	FI	Solid
s_1_315	GE	Solid
s_2_315	GF	Solid

This data looks very clean, noting that 2 were routed through NSV, thus changing the .txt file column. The biggest thing of note is that the TCs on hp_45 indicate that heat pipe was not transferring power quickly during startup (likely had limited throughput at low temperatures more than the other heat pipes), which matches data for the hp_1_45 and hp_2_45 TCs, so the readings are likely true.

Simulator Body Thermocouples

These thermocouples are welded to the outside of the Stirling simulators. If you have read the paper on the system test results², then you know that this is where things start to get interesting, in terms of determining the power lost from the simulators. The locations of these TCs be seen on the left photo in Figure 5, and the locations are indicated on Figure 4 as well. The s_3 TCs are on the narrow neck portion of the simulators, and the s_4 TCs are approximately halfway up on the larger diameter section of the simulator.

TC name	.txt file column	Trusty? based on full run data
s_3_0	BX	Solid
s_4_0	BY	Solid
s_3_45		Gone
s_4_45		Gone
s_3_135	DU	Solid

s_4_135	DV	Solid
s_3_180	EJ	Solid
s_4_180		Gone
s_3_225	FA	Solid
s_4_225	FB	Solid
s_3_315	GG	Solid
s_4_315	GH	Solid

Modeling heat loss from the simulators is tricky because the top portion was uninsulated. About 2" below the top of the simulator is uninsulated, plus the top and the flange. The size of the uninsulated portion under the flange varied between .5" and 1", as can be seen on the right photo in Figure 5. During nominal operation this was not a big deal because the gas-flow kept the top and outside of the simulators cold. But when simulator flow was cut the top of the simulators appeared to radiate a significant amount of heat.




Figure 5. Simulator and environmental TC locations

Environmental and Sink Thermocouples

Several TCs were positioned to measure the environment or heat sink next the convertors and simulators. Four TCs (al_1) were mounted the cylindrical aluminum structure located inside of the convertors and simulators, and another 4 TCs (es_1) were left floating (unbonded) near the structure. The arrows pointing to the aluminum structure in Figure 5 are pointing to possible locations for these TCs.

TC name	Location	.txt file column	Trusty? based on full run data
al_1_0	Fastened to aluminum structure a 22.5°	CB	Solid
es-1_0	Free floating near aluminum structure near 22.5°	CC	Solid
al_1_45	Fastened to aluminum structure a 67.5°	CS	Solid
es_1_45	Free floating near aluminum structure at 67.5°	CT	Squirrelly
al_1_135	Fastened to aluminum structure a 157.5°	DY	Solid
es_1_135	Free floating near aluminum structure near 157.5°	DZ	Solid
al_1_225	Fastened to aluminum structure a 247.5°	FE	Solid
es_1_225	Free floating near aluminum structure at 247.5°	FF	Squirrelly
mli top	Taped to MLI surrounding collector	CK	Solid
mli middle	Taped to MLI surrounding HP/simulator interface	CL	Solid
mli bottom	Taped to MLI surrounding condenser	CO	Solid
mli heatpipe	Taped to MLI surrounding HP just below condenser	CP	Solid

The Al structure served as the anchor for brackets that supported each engine and simulator, this is best seen in Figure 6. The Al structure also provided a central region to protect/isolate the gas inlet and outlet flow tubes. The change in structure temperature is therefore coupled to the outlet flow temperatures, and the inlet flow temperatures are subsequently coupled to the structure temperature. The effective sink TCs are similar to the Al structure TCs, but not as coupled to the flows, and more sensitive to the aforementioned heat loss from the uninsulated top of the simulators (most noticeable when flow is turned off).

One additional note is that the mli top TC (shown in Figure 5) ended up being covered up by the MLI that surrounded the simulator body; notice how the large cylinder of MLI truncates the top/collector cone on the right photo. This likely caused mli top to read higher than it otherwise would have.

Gas Inlet and Outlet Thermocouples

The thermocouples that measured N2 gas inlet and outlet temperatures are shown below. The flow rate for the Stirling engines was very high and continued throughout the test to keep the “cold end” cold, as if it was attached to a cold radiator (although the N2 flow was likely 50 to 100 C colder than what a practical space radiator might deliver). For the simulators, the gas inlet temperature (s_5) is measured just before the gas enters the top of the simulator. The gas outlet temperature (s_6) is measured inside the simulators with a TC that penetrates into the simulator along the radial C/L: this TC is estimated to be 1” above the copper heat acceptor at the bottom of the simulator. The locations of these TCs can be

seen on Figure 6: the outlet-TC penetrates from the top, and the inlet-TC is exiting horizontally to the left of the simulator C/L.

TC name	.txt file column	Trusty? based on full run data
s_5_0	HU	Solid
s_6_0	HV	Solid
s_5_45	BP	Solid
s_6_45	HY	Solid
e1_9_90	DI	Solid
e1_10_90	IE	Solid
s_5_135	EC	Squirrely
s_6_135	IH	Solid
s_5_180	FY	Solid
s_6_180	IJ	Solid
s_5_225	BO	Solid
s_6_225	IM	Solid
e2_9_270	IR	Solid
e2_10_270	IS	Solid
s_5_315	ET	Solid
s_6_315	IU	Solid

Only the s_5_135 TC appeared to have any trouble. The readings become suspiciously higher between $\sim T=2$ to $\sim T=4.5$ hours. The other TCs tracked pretty well with each other. The inlet gas temperature for the simulators increased significantly as the “guts” of the PCS warmed up, from ~ 15 C to ~ 90 C. After that, the inlet temperature was highly dependent on flow rate, ranging from ~ 60 C for the very high flow rate cases, ~ 90 C at nominal flow rate, and ~ 200 C with no flow rate (perhaps coming into equilibrium with the heat sink environment, or the effective temperature of the guts). For the outlet temperature, some of the changes that occurred during supposedly constant flow-rate periods give the first hint that the flow meters were acting squirrely. The other thing of note is that the outlet flow TC went to 770 C when there was no flow, which is very close to the collector temperature, and it is likely that the bulk of the gas inside the simulator approached this temperature (via conduction or a small amount of internal natural convection). If the hot simulator gas worked its way to the top of the simulator (which seems more likely than not), this would cause rather large passive losses from the top of the uninsulated simulator section.




Figure 6. Gas routing and thermocouples

FLOW RATES AND MISC.

Stirling Convertor and Simulator Flow Rates

The flow rates of N₂ through the convertors and simulators were recorded in standard liters per minute (SLPM); i.e. liter a standard temperature and pressure. This corresponds to .0446 moles/min, and for N₂ 1 SLPM = .0208 g/s. The power draw by the simulators can be calculated using the flow rate times the change in gas temperature times the specific heat. The flow rate though the convertors remains high throughout the entire full-run, to keep the alternator cool, although there is very little heat transfer to the cold end if the piston is not moving. The simulator flow travels down through the outside of the simulator, reverses direction at the bottom, and flows up through a copper HX that mimics the acceptor of the convertors, and then up and out through the center of the simulator. In the large diameter portion of the simulator (above the copper HX), the downward inlet stream and upward outlet stream are separated by a gap that contains 4-wraps of MLI.

During the first several hours of the full-run, the calculated power draw seems to agree very well the experimental data, but there are times during the full-run where there are significant differences (10s of percent) between modeled and actual power. The data, and anecdotes from personnel, indicate that there were problems with the flowmeter readings. This is alluded to in the results paper, but I will give more detail here. Not only do many of the flow rates read negative at “zero-flow”, the value of this zero-point changed significantly throughout each of the tests, most notably the full-power run. This indicates the calibration did not hold – either because of changes in the zero-point and/or changes in the expected gain (or multiplying constant) per unit of flow. The zero-point reading (i.e. when the flow was set to zero) at various points in the transient are listed in the table below. There is some correlation among the individual readings with respect to each other. There also could be some correlation with ambient temperature near the flowmeters or perhaps the power level (flux, Rad environment), but this is not evident. Calibration might have been caused by changes in the power signature provided to the meters, whether it was random or induced by changes in electrical power draw throughout the components of the experiment and facility.

It is also important to note that when I talked to Marc Gibson about this issue, he recalled that attempting to calibrate the flow meters in the DAF was very difficult and odd; whereas it was much more easier/straightforward at GRC. So the issue could have been related to facility power/ground, instead of, or in addition to the flowmeters themselves (unless the flow meters were damaged in transport from Ohio to Nevada, but it would be odd for all of them to behave strangely, so it is more likely the power supply).

The zero-point reading (i.e. when the flow was set to zero) at various points in the transient are listed in the table below. There is some correlations among the individual readings with respect to each other. There also could be some correlation with the environment temperature of the flowmeters and perhaps the reactor power level (flux, Rad environment).

Flowmeter name	.txt file column	15 cent run T<~2	15 cent run T>~2	30 cent run T<~1.5	30 cent run T>~1.8	60 cent run T<~5.5	60 cent run T<~6.0	Full run T=0.0	Zero-point T=12.0	Zero-point T=13.0	Zero-point T=14.0	Zero-point T=22.7	Zero-point T=25.0	Zero-point T=27.1	Zero-point T=29.0	Zero-point T=72.0
s_1_0	IY	-2.1	-9.4	-9.9	-2.0	-9.4	-1.8	-8.7	0.5	0.6	0.7	-8.9	-13.2	-9.3	-9.7	-4.3
s_1_45	IZ	-1.9	-8.4	-8.9	-1.9	-8.4	-1.6	-7.7				-3.9	-9.0	-4.5	-4.3	0.0
e_1_90	JA	-1.3	-0.4	-1.7	-1.3	-1.7	-1.8	-1.7								3.2
s_1_135	JB	-2.4	-11.3	-11.7	-2.3	-11.2	-2.1	-10.5				-7.2	-11.3	-8.9	-9.5	-3.0
s_1_180	JC	-2.4	-11.3	-11.7	-2.4	-11.3	-2.2	-10.5		0.5	0.1	-6.0	-8.9	-6.0	-6.4	-0.3
s_1_225	JD	-2.2	-11.1	-11.5	-2.1	-11.1	-1.9	-10.3				-5.5	-8.3	-4.0	-4.0	1.1
e_1_270	JE	-1.8	-0.8	-2.3	-1.7	-2.4	-2.6	-2.7								2.2
s_1_315	JF	-2.3	-11.2	-11.6	-2.2	-11.1	-2.0	-10.4				-4.4	-7.3	-3.2	-3.4	1.7

The data for the free runs, when there was no N2 flow (except for an early test in the 60 cent run), each show a dramatic ~step change in zero-point during each test. The 15-cent run starts with the zero-points looking relatively good (~-2.0), but then has a rapid transition to much lower zero-points about 2 hours in. The 30-cent and 60-cent runs start with similar low-zero points, and then transition back to

“better” zero-points during the transient (but the subsequent transient revert to starting with the low zero-points). The full-run again starts with the strongly negative zero-points, then two of the simulator flow meters read close to zero later on, then readings vary considerably after that. Finally, while the system “rests” and passively cools over the weekend, the zero-points all get closer to zero.

One final thing to note is that the convertor flowmeters do not vary nearly as much as the simulator flowmeters. This is likely because the convertor flowmeters were indeed different models; they were flowmeters rated to much higher flow levels, so therefore maybe less sensitive to changes in input power signature (voltage, ground, stability, etc.).

Overall, these flow readings, in combination of the gas-in and gas-out TCs are obviously very important to attempt to accurately match the power output of the reactor. My hope is that it may be possible to figure out how the flowmeters were behaving, and come up with new set of estimated flow rates that match the data (which still have some connection to the actual data). I was looking at this when I was forced to move onto other projects, but hope to find more time in the future. It appears likely that the flowmeters were reading too high (versus the actual flow) during the high power transients, in particular the 2nd one. The readings for the zero-points at $T \sim 12$ indeed indicate that major calibration changes could have occurred during the time frame of the high power transients.

While it’s important to examine the negatives associated with this flowmeter issue, it does not significantly change the incredibly positive, confirmatory value of the KRUSTY test. The overall results of the test, including the 60-cent run, the early transients, and the Stirling restart transient, provided solid data that verify that the core thermal power matched the power loss from the fuel (which is in agreement with the simple physics); the flowmeter issue simply prevented the potential for additional confirmatory examples. Also, even though the flowmeters were potentially 10s of percent uncertain, the higher power transients indeed verified successful reactor operation at higher fission powers

Vessel Bleed Flow Rate

The bleed flow, or vacuum-can flow, is discussed in the Nuclear Technology “results” paper. This output is in .txt file column JG and labeled “nsv_can_1_nitrogen_flow”. According to NASA, the number printed is percent of full flow, and there was no attempt to calibrate to cfm. Regardless, the use of this flow reading, other than via a normalized heat removal factor, is probably not going to be practical. Also, note that the zero-point is -10.5 until it is turned on. Interesting that it reads similar to most of the simulator flow rates.

Vacuum Level

This reading might have significance for a serious dynamic benchmarking effort. The output is in .txt file column IV, labeled as “nsv_vac_1_vacuum_level”, and is in the units of Torr. As noted in the Nuclear Technology “warm crits” paper, the vacuum turbopump died near the end of the 60-cent run, and the reading went from a solid $1e-4$ Torr to $\sim 1e-2$ Torr. For the full-run, the reading starts at $\sim .006$ Torr and gradually increases as the system temperature increases. A lot of the initial increase might have been due to warming up of the chamber gas. Then there is a significant spike in pressure when the system is heated to its highest level at $T \sim 18.3$. This temperature was higher than the core was heated to in the electrically-heated checkout prior to nuclear testing, thus there was likely a significant amount of off-gassing. These changes in pressure in the 60-cent run affected heat transfer slightly (it appeared to

change the oscillation shape after pump failure in the 60-cent run). In the final-run, the pressure got high enough that gas conduction may have been significant, and also some of the surfaces were significantly blackened due to the combination of high temperature and significant gas pressure (of unknown composition). Figure 7 shows how blackened the outer surface of the fuel was after final disassembly.



Pre-test fuel photo



Post-test fuel photo

Figure 7. Pre- and post-test photos of fuel piece

Stirling Convertor Parameters

There are numerous column in the .txt file that pertain to the condition of the Stirling convertors. While electrical production is one of the important outcomes of the test, it is not relevant to system dynamic benchmarking. The electrical output of the engines themselves is a function of commanded voltage, hot-end, and cold-end temperatures, which was studied extensively by NASA during stand-alone and electrical system testing. Any benchmarking of engine electrical performance should be focus on stand-alone convertor performance data from NASA. A chart showing the electrical power and power balance throughout the entire full-run in shown in Figure 11 of Reference #4. The best numbers to use for electrical power are in .txt column AC (AC = A/C = coincidence?) e1_altpower_estimate3_w_calc and .txt column BA e2_altpower_estimate3_w_calc. This number is in Watts, and is an estimate based on reliable variables that was deemed by NASA to be very accurate. The convertors were rigorously tested by NASA as stand-alone convertors and in the system configuration via electrically heated testing, and this predicted value matched very closely to actual output.

The parameter of concern for system dynamic benchmarking is thermal power draw of the convertors, which is also a function of commanded voltage and temperatures. As a benchmarking cheat, a modeler could pull the "thermal power in" data off of Figure 11⁴ to set the power draw, but since power draw is a strong function of the temperature, a consistent model would calculate the temperature and base the power draw on the engine characteristics. Fortunately, the cold end temperature throughout KRUSTY was relatively constant (~15 C), so a simple 2-variable correlation for power draw could be developed based on the stand-alone NASA testing; based on hot end (collector) temperature and the commanded voltage to the convertor (which determines the engine stroke).

POST-TEST INDICATIONS OF THERMAL BONDING

Photos taken during KRUSTY disassembly, plus anecdotes from the engineers that performed the disassembly, can provide several clues to how well components were thermally bonded to each other. Figure 8 provides indications of how well each fuel piece was thermal bonded in the axial direction. The photo on the right shows the fuel column prior to full disassembly, and indicates that there was no appreciable gap between the fuel pieces. There was some concern that after the first heating to 800 C (the electrical checkout test) a gap would be created between the fuel pieces as the core cooled (due to differential expansion and slippage between the heat pipes and the fuel itself during heat up). This could have lessened the feedback during startup of each subsequent nuclear test (i.e. as fuel warmed and expanded it would close the potential gap). Fortunately this photo indicates that such a cold gap never occurred.

At 800 C, it was always anticipated that each of the fuel pieces would be in contact, but uncertain if there would be tight contact. The photos seem to indicate that hot gas was able to infiltrate the gap on the outer and inner edges, but not in the middle. The bottom and middle pieces appeared to have good contact over the majority of the middle section, whereas the middle and top pieces had to have a stronger contact over a smaller area (only $\sim 1/2''$ width). A very detailed model might vary fuel-to-fuel conductivity as a function of radius, but it is unlikely to change results significantly (because there is no benchmark data for internal fuel temperatures, only the edge). It was noted that the fuel pieces came apart easily from each other during disassembly, and thus did not physically bond together.




Figure 8. Post-test photos of fuel

Post-test photos can also give an impression of axial conduction out of the core. The left photo in Figure 9 shows that the upper reflector and the SS and B4C plug shields were well aligned with each other, and likely in some contact. These components were resting upon each other, so the compression force should be a function of the weight of the plug shields. The orange disc in this photo is the vessel O-ring. The middle photo shows the MLI after the upper fuel piece was removed. The MLI shows a significant discolored region directly above the central hole in the fuel (which is where the hot fuel outgassing would have attacked). Note that the small hole in the upper MLI and BeO was for the top of the neutron source to be aligned in early zero-power criticals. Two features indicate that the upper fuel was in at least modest compression with the upper MLI; a lack of discolorization and an apparent physical imprint where the fuel would have aligned with the MLI. The expected force/pressure at this interface would be the weight of the upper pieces (the BeO and the plug shields). The MLI appeared to be have about the same foil separation as before the test, such that the threads were not permanently crushed. The right picture shows significant darkening of the BeO near the hole in the MLI, but the surface is relatively clean elsewhere (but surely of emissivity of the entire piece was higher than before the test).

The other interesting feature of the mid and right photos is the copper attached to the heat pipes. The heat and pressure of the test caused a physical bond (although rather loose) between the copper and the Haynes 230 heat pipes; noting that the copper did not bond significantly to the fuel. In the right photo it appears that on the rear heat pipe, some of the copper is gone in the azimuthal-middle where the force would have been highest; perhaps the copper was squeezed to the sides? Other photos show the same sort of profile in the copper.





Figure 9. Post-test photos upper axial region

Figure 10 shows the lower axial reflector and MLI. In this case the MLI was physically more compacted than the upper MLI; this was expected because, as discussed earlier, this is the interface where the downward thermal expansion of the core conflicted with the lesser downward thermal expansion of the vacuum-can. This MLI is relatively clean, and a zoomed view of the photo actually shows the thermocouple wires ab_90 and ab_270 (although they are now free floating, so it is hard to know where the junction resided when the fuel was in place – it would depend on the gaps/forces as the wire was inserted).

The most dramatic feature of Figure 10 is the remarkably clean fuel surface where it interfaced with the shield plug (as was also indicated in the fuel photo above). The size of this ring is a mirror image of the off-white outer ring on the upper shield plug; whereas the middle of the plug has been severely blackened by the hot outgassing just as the upper mli was. These photos answer the question of how heat escaped the bottom of the core. The central BeO plug was apparently in very tight contact with the fuel (there was no MLI here). Also, the same force would have been thermally bonding the upper shield plug to the lower shield plug, and the lower shield plug to the boss on the vessel end cap. Thus, there was likely a very good conduction path from the fuel to the bottom of the vacuum can over the area of the plugs. There was also probably pretty good conduction from the fuel to the outer BeO piece, but at a much lower gap conductance. The other thing of note if the photo on the right of Figure 10 are some of the mli pieces that were wrapped around the bottoms of the heat pipes. These wraps came off the heat pipes as the core was pulled out the vacuum-can. There is some worry that one or more of these pieces could have slid down partially off the heat pipe during installation or testing, which would have exposed a portion of the heat pipe just below the fuel. There is no basis for assuming that this happened, but it is a possibility; such that more heat would be lost from the heat pipes to the BeO and the radial MLI in this region (the bottom in the left photo of Figure 2 shows this MLI in place prior to the test).



VARIOUS THOUGHTS

The Relative Importance of the Neutron Source

An AmBe neutron source was in the axial-center radial shield port for all of the nuclear-powered tests; this location is where the lower of the 2 SS plug cylinders protrudes from the left side of the shield on the left photo of Figure 3. An email from Rene Sanchez quoted the source strength at 1.13×10^6 with dimensions of .69" D x .76" L. The source was placed in a holder. More details are in Reference 5.

A detailed neutron source model is not generally needed for dynamic modeling, unless an attempt is made to recreate startup with the physical neutron source in place, versus a generic source term that creates the power level during the 15-cent ramp. The latter seems like cheating, i.e. creating the source term based on the results, but remember that we used results to determine that we had a ~15-cent insertion in the first place. The level of reactivity insertion is based on tracking the actual data, not by determining the precise position of the platen. The ultimate end-to-end transient model would model the neutron source explicitly, use integrated 3D transport based neutronics, and track flux/power based on platen position and system temperature/geometry changes. I hope someone in the future has that capability, but we are long way from practically doing that in 2020 (especially for the full-run).

On the subject of source, there is one effect that plays a minor academic role – photoneutrons. The impact of photoneutrons on reactivity insignificant, $\ll 1$ cent in keff. However, the effect of photoneutrons appears to limit the power drop after reactivity insertion is complete in the warm criticals. When I first modeled the warm criticals, I was not matching the drop in power/flux after reactivity insertion was complete. My calculated power was significantly low, albeit only noticeable on a log scale. In that scenario, the reactor is subcritical and the “bottoming” flux is based on the multiplying the neutron source. I then realized that photoneutrons from decay gammas were probably a stronger than the effect of the AmBe source. I added a term for a photoneutron source in FRINK that was a linear function of decay power (hoping that the high energy gamma source was relatively close to power source). Through empirical normalization I was able to use this source term to match the low power trajectory of each case – the 60-cent run is the most noticeable because it had the highest decay power of the three. This effect for the full-power run after scram is probably not evident (despite the higher decay power, because keff quickly drops to a point where neutron multiplication is very small).

The Importance of Platen Position

If you’re are performing criticality benchmarks, then clearly the platen position must be known for the specific configuration you are evaluating. If you are benchmarking the warm criticals (or perhaps a load following transient where the platen doesn’t move), then the platen position is not necessarily important, unless your system model includes neutronic calculations, instead of kinetics (which incorporate reactivity feedback correlations). A kinetics-based model can simply insert ~15 cents of reactivity at the start, and then continuously add reactivity to reach 30 cents and 60 cents during the time period shown in the experimental data. Keep in mind that the insertion numbers are not exactly 15, 30 and 60 cents – the indicated initial ramp reactivity is on Table 1 above. The first test actually inserted closer to 15.5 cents, so we stuck with that approximate level throughout to keep the start the same for all cases. There is no “correct” answer for the final reactivity insertion for the 60-cent run, because we estimated where the platen should be to achieve that number – the actual data indicate that less the 60 cents was inserted.

Corrections to ANS Nuclear Technology Papers

The design paper³ states that the lower axial reflector is 4 pieces, but there were only 3. The internal plug was split into two 2” tall pieces, but the outer annulus was one 4” tall piece. The paper may also imply that MLI was placed between the fuel and entire lower axial reflector, but MLI was only placed on top of the outer/annular BeO piece, not over the plug piece (as described above in Figure 10). It was actually intended that insulation be attached to the top of the upper plug piece, but it appears that it was not added.

Thoughts on Heat Pipe Modeling

Physics-based transient heat-pipe modeling is extremely difficult, and perhaps impossible to practically implement within a system-wide reactor dynamics model. One of the positive attributes of heat pipe reactors is that when the HPs are operating well within their throughput limits, the internal/vapor region of the HP passively acts as a quasi-infinite conductor. Thus, there are numerous options for modeling a heat pipe in a transient reactor code; ranging from an infinite conductor to a full thermal-hydraulic, compressible flow, 2-phase CFD model (good luck on getting that to work within a transient reactor model!).

The easiest option is to model the HP as an infinite axial conductor; either with an insanely high thermal conductivity, or direct coupling between the evaporator and condenser nodes. A decent model will at least model the HP wall as separate from the internals (and potentially the wick, annular and/or film regions separate from the vapor), where the “infinite” conductivity is applied to a singular central node (approximating the vapor region) within the HP. It is important to ensure that most of the HP thermal capacitance is not skipped in this process. Of course, the internal conductance is not infinite (i.e. zero delta-T) because there is a pressure drop along the HP which causes the vapor to cool (via expansion) slightly as it flows downstream, plus it can lose heat to the cooler wick/film as it flows downstream and pick up cooler entrained vapor. In the end, what matters is that the saturation temperature is lower in the condenser than the evaporator. The amount of internal delta-T depends on the heat pipe design and temperature. A detailed heat pipe model, i.e. one that models the internal flow, might do a good job estimating this delta-T, but the alternative is to use whatever steady-state test data exists to come up with an empirical steady-state thermal conductivity. Usually all temperature measurements will be on the outer HP wall (except for HPs specifically made for research), so the radial delta-Ts (which are often larger than the internal, saturation-temperature delta-T) must be subtracted to estimate an internal delta-T, which can then be used to estimate a thermal conductivity for the internal node. For KRUSTY, the conductivity value I have used has varied as more data came in over the course of the program – the current value which I use to represent the “unlimited” internal conductivity is 4 MW/m-K (sodium HPs are truly astounding at moving heat!).

The major caveat to all of the above is “when the HP is operating well within its throughput limits”, which is never true (by definition) during startup. The good news is that to model most KRUSTY transients, the high internal conductivity method should work very well, except for startup and shutdown. Also, for the warm criticals the internal heat pipe conductivity can be set to zero (unless you want to try and get extremely accurate on the 60-cent run). So, the only time that a better model is really needed is for startup of the full-power run.

The heat-pipe modeling technique I use in FRINK is to use the steady-state HP throughput limits to modify/fudge the internal heat-pipe conductivity as needed, i.e. if one of the limits is exceeded at startup, shutdown, and any time during operation. The first step in this process is to come up with a set of throughput limits; either via a steady-state model or empirically via test data (which is preferred if you can get the data). As mentioned in each of the Nuclear Technology references, the KRUSTY HPs only have a wick in the pool/evaporator region, and use thermosiphon action (gravity) in the rest of the “heat pipe” (purists can debate whether they should be called heat pipes, thermosiphons, or some sort of hybrid name). Heat pipes and thermosiphons have physical limits to throughput: for KRUSTY the key limits are the sonic limit (speed of sound of vapor), viscous limit (how quickly the Na can flow down the

un-wicked region due to gravity, and to a lesser extent how the wick distributes flow in the evaporator) and flooding limit (where liquid down-flow restricts the area for vapor up-flow and/or interacts via mixing/entrainment – the area restriction and mixing are drawbacks of a thermosiphon vs heat pipe (which controls where/how the fluid flows), while HPs still can still have entrainment limits).

The throughput limits determined by NASA (via modeling and testing) for the KRUSTY heat pipes are shown in Figure 7 of Reference #2; that figure also shows when it appears the limits were hit during the testing. Within FRINK, I check whether the heat transfer (watts) being calculated along the axial central node exceeds any of these limit values. If it doesn't, the node conductivity stays at its quasi-infinite value; which in theory should change with temperature and throughput, but I just use a constant value, currently the 4 MW/m-k mentioned above. If the calculated axial conduction exceeds one or more of the limits, then I proportionally reduce the conductivity by the ratio of the calculated Q to the worst-case limit; i.e. if the viscous limit is 10 W but I am calculating 100 W with the unlimited conductivity value (based on the delta-T of the node from the prior iteration), I then reduce the conductivity of that node by a factor of 10. If the HP is below 400 C, the limit is ~zero so this ratio sets the axial conductivity to ~zero.

I use this throughput-limit-ratio method on an axial node-by-node basis, but it gets even messier from there. I treat the sonic and flooding limits as local limits, based on the temperature of the internal radial HP nodes adjacent to the central/fudged vapor region. In effect, if the wall is subcooled in the next downstream HP node, the limits will decrease and heat transfer is impeded. In my modeling thus far, this has done a good job of recreating the thaw front as the heat pipe starts. At startup, downstream nodes are at room temperature, so the axial heat transfer halts (apart from material conduction) when heat reaches the first significantly subcooled node (initially just outside the core), until that node heats up enough to increase throughput limits enough to get heat to the next node, and so on. However, the local application of the HP limits may not work very well with the viscous limit, which should in theory depend on how far a node is from the evaporator (note: all of the HP throughput limits are integral steady-state limits for the entire heat pipe, but it seems logical that the sonic and flooding could be dictated more by local temperature and radial dimensions, while the viscous limit would depend on integral temperature and the required flow-length to reach that node. To implement the viscous limit I apply a small fraction of the viscous limit at every node, and then add the remaining fraction of the limit based on how far the node is from the evaporator (actually the ratio of the node's distance over the total distance from evaporator to condenser); therefore, to fully reach the condenser the entire viscous limit must not be exceeded (and if it is exceeded, Q (conductivity) will be reduced accordingly).

This method is at best a decent quasi-steady approximation to what is actually going on inside the HP, and at worst, not as good as sticking with the fixed conductivity model; but I think it makes some sense, and it was the simplest thing I could think of to implement and not substantially slow down FRINK execution time. This model probably fits the mantra "something that appears reasonable is better than nothing", rather than "best is the enemy of good enough" – at least it seems to do reasonably well during KRUSTY startup.

Another aspect of transient heat pipe modeling is the pool, specifically the depletion of the pool as the heat pipes start to transport heat. The pool height affects reactivity modestly, because when operating much of the Na lost from the pool leaves the core. In FRINK I estimate the size of the pool based on power throughput (actually power history over a short time-frame), and apply the pool height to a

reactivity worth curve I generated with MCNP. This effect is discussed in Reference #2, and is noticeable in the experimental data, but it can be left out of simpler models without much consequence. This calculation also implies a dryout limit, if the pool and wick are depleted of Na, which would occur in an under-filled HP

Other Suggestions for Dynamic Modelers

The first suggestion is obvious. Start benchmarking with a course mesh and simplified model. None of the warm criticals have azimuthal asymmetry, so only a 1/16th core model is needed to include the effect of the heat pipe and conduction heat transfer to the core ring-clamps. In fact, the heat pipes themselves do not play a major role in these transients (except a minor role in the 60-cent run), so you could start with a model that has no azimuthal nodes. Perhaps no axial nodes as well and only model radial heat loss initially. The most important thing early on is to have a model that runs quickly, so you don't have to wait hours to find out if/what adjustments you need to make.

The worst thing a modeler can do is stress about nailing down details early in the benchmarking process. When you are unsure, make guesses as best you can, so you can start to make progress; actually when you are making guesses you will learn more about how the reactor operates and gain intuition on how one thing affects another. Then you may find that the parameter you were stressing about doesn't make much difference, or if it does, you now know it is definitely worth the time to dig deeper. Also, if you have a graph of a parameter you need, but can't find the correlation, don't wait until if/when you find it. Read x,y values off the chart and then interpolate within your model or create your own curve fit correlation – if you never do this or other simplifications/guesses, you're probably not moving fast enough in your modeling effort.

If a detailed model is attempted, then many guesses/fudges will have to be made to try and match the data. One example might be changes in material emissivities. Post-test photos clearly show a highly blackened fuel. Post-test photos also show some darkening of the MLI, especially the 0 and 315 simulators, so a more accurate model might fudge increases in emissivity for various components (maybe as a function of temperature and chamber pressure history), and see if that matches the data better. Also, the increase in chamber pressure as the test progressed may have added a conduction loss component, which would also increase passive thermal losses as time progressed. Similarly, a model that increases gap conductance between the fuel/HPs/rings versus temperature, or an axial MLI conductance correlation due to compression will probably fit the data better than a constant model. The "nominal" correlation of MLI conductance is another tricky topic that requires guesswork (e.g. how what fraction of the threads/foil are in contact). I wrote a separate code to develop MLI conductance correlations as a function of hot-side and cold-side temperature, based on several input assumptions, until I got one that seemed to match fairly well.

If done well, then all of your guesses/fudges can ultimately represent valid empirical relationships. Even better if most aspects of these guesses are tied to data and/or based on physics (i.e. a function of temperature, pressure, time-history, etc.). All this said, it is very easy to go down the rabbit hole and convince yourself that you've done a good job with your guesses, while you might have just gotten lucky that your guesses/fudges offset each other and appear to create the right answer, but won't be correct for future use.

Don't forget decay power! The decay power changes throughout the transient, so you'll need to account for it if you want to improve your accuracy. The decay power should effectively vary between 3% and 6% throughout the transients; the timescale of the first several percent is so short that it is quickly saturated, and the reactor does not operate long enough to get close to 7%. The neutron detectors are measuring the neutron flux, thus indicating the fission power, not total thermal power (except for the small amount of delayed neutrons). A good transient model will track decay power based on the time-history of the transients – I use a decay-power precursor technique similar to how delayed neutrons are tracked.

Another suggestion is to frequently run the code with major simplifications for sanity checks. Perhaps run cases with all heat transfer from the fuel set to zero (or just the axial heat transfer, etc.), and see if the deviation from the reference case makes sense. Similarly run cases with 2x your reactivity feedback, or conduction, and especially your fudge factors to see if the results make sense (and if you don't know if they make sense, you might start to develop an intuition for what makes sense). Perhaps try running with constant properties/feedback initially, instead of temperature dependent values, and leave the ability to check the difference in the future, to see if the results make sense. This is how I've developed FRINK (which at least for KRUSTY did very well), and more importantly how I've gained intuition into how reactors operate and how to better design them.

The data files are pretty large given a 1 second interval. Consider reducing the data to 10 second intervals so that when you plot it and compare it, your output files and excel files (or whatever you use to plot/compare) are of manageable size.

Finally, frequently interrogate the calculations. Hopefully, the code you are using will have the ability to check the value of any parameter at any time (I just add print statements wherever, whenever I need them, but if you're using a fixed-compiled code you have to do as best as you can). If you can't check any specific calculation or section of code, find ways to perform simple hand/excel calcs as sanity checks, which goes with the previous paragraph. Fortunately, these warm criticals provide pretty simple benchmarks, but there is still a chance to end up with matching results via a fortuitous combination of mistakes. This happens a lot in complex modeling efforts because the first time you get the result you expect you assume everything is correct and declare victory.

Do not be afraid to take shortcuts to get something working. Just as for KRUSTY – best is the enemy of good enough.

CLOSING REMARKS

This document is largely intended to aid future benchmarkers of the KRUSTY tests. I'm trying to document every fact and opinion I can remember about the testing, which should been done earlier, but better late than never; although I don't think anything significant has been lost or forgotten that would have been remembered 2 years ago. The majority of the information provided is probably only useful for those who attempt to model the entire system, including the power conversion system. There are likely specific details that are not available to account for every physical effect, but in those cases an educated guess will hopefully be sufficient. The purpose was to publish every bit of raw data, and then give context to that data with recollections and opinions, especially in cases where instrumentation appears questionable.

The difficulty of obtaining accurate and dependable instrumentation for KRUSTY should be viewed as a warning to future reactor programs. Fortunately, KRUSTY was adequately instrumented to achieve all important goals regardless of the aforementioned instrumentation issues (plus a single-event-upset trip of one of the rack computers). Success was achieved because KRUSTY was a very simple reactor which did not require real time reactor control (except for the COMET scram system). Any concept that relies on reliable instrumentation to control the reactor in real time, must allow significant time early in reactor testing to diagnose, interpret, and potentially fix instrumentation issues. This is repeatedly mentioned in documentation of the Rover/NERVA tests; they spend a lot of test iterations determining how to instrument the next test, and for each individual test they performed a lot of preliminary state-points and transients to interpret which instruments appeared accurate and which didn't (like we've done for KRUSTY, but only after the fact). The biggest implication of this is for anyone considering a flight demonstration of a space reactor; if the reactor doesn't have simple physics and control, then a flight demo carries a lot of risk. Failure may be likely for a reactor that depends on accurate/reliable instrumentation for control and/or there is no time to perform numerous low-power/temperature operations to interpret which readings can be trusted more than others.

Personally, I would not necessarily view the difficulties we experienced in instrumentation, or our lack of rigor in recording decisions/details, as a reason for strict QA in engineering, technology, and testing. Any reactor project must focus on what areas require the most attention to meet big picture goals (e.g. complete a test) versus less important goals (e.g. ensure every aspect of the test is highly prototypic and accurately recorded). Whenever possible, KRUSTY used the "best is the enemy of good enough" philosophy, which ultimately may have been the difference between success and failure (i.e. all novel reactor programs in the US over the past 40 years).

REFERENCES

1. Poston, D. I., Gibson, M. A., McClure, P. R., & Sanchez, R. G., "Results of the KRUSTY Warm Critical Experiments". *Nuclear Technology*, 206(sup1), 78-88 (2020).
2. Poston, D. I., Gibson, M. A., Sanchez, R. G., & McClure, P. R., "Results of the KRUSTY nuclear system test". *Nuclear Technology*, 206(sup1), 89-117 (2020).
3. Poston, D. I., Gibson, M. A., Godfroy, T., & McClure, P. R., "KRUSTY Reactor Design". *Nuclear Technology*, 206(sup1), 13-30 (2020).
4. Gibson, M. A., Poston, D. I., McClure, P. R., Sanzi, J. L., Godfroy, T. J., Briggs, M. H., ... & Lugasy, N., "Heat Transport and Power Conversion of the Kilowatt Reactor Test. *Nuclear Technology*, 206(sup1), 31-42 (2020).
5. Sanchez, R., Grove, T., Hayes, D., Goda, J., McKenzie, G., Hutchinson, J., ... & Smith, K., "Kilowatt Reactor Using Stirling Technology (KRUSTY) Component-Critical Experiments". *Nuclear Technology*, 206(sup1), 56-67 (2020).
6. Grove, T., Hayes, D., Goda, J., McKenzie, G., Hutchinson, J., Cutler, T., ... & Sanchez, R., "Kilowatt Reactor Using Stirling Technology (KRUSTY) Cold Critical Measurements". *Nuclear Technology*, 206(sup1), 68-77 (2020).

7. McClure, P. R., Poston, D. I., Gibson, M. A., Mason, L. S., & Robinson, R. C., "Kilopower Project: The KRUSTY Fission Power Experiment and Potential Missions". *Nuclear Technology*, 206(sup1), 1-12 (2020).
8. Smith, K. N., et al., "KRUSTY: Beryllium-Oxide and Stainless-Steel Reflected Cylinder of HEU Metal," HEU-METFAST-101, Nuclear Energy Agency/OECD (Organization for Economic Co-operation and Development) (October 2019).